

REPORT DOCUMENTATION PAGE

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE (DD-MM-YYYY) 15-01-2009		2. REPORT TYPE final		3. DATES COVERED (From - To) Feb 2005 - Dec 2008	
4. TITLE AND SUBTITLE MODELING FOR SEMI-AUTONOMOUS CONTROL OF COOPERATIVE VEHICLES				5a. CONTRACT NUMBER FA9550-05-1-0118	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Campbell, Mark, E				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Cornell University 208 Upson Hall Ithaca NY 14853				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NL 875 N Randolph St Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution A: Approved for Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report details work completed in an AFOSR program focused on semi-autonomous control of cooperative vehicles. Coupled operator-multiple vehicle systems are modeled in a unified framework using probabilistic graphs to yield a methodology for analyzing semi-autonomous systems. The framework uses conditional probabilistic dependencies between elements, leading to a Bayesian network with probabilistic evaluation capability. Both discrete and continuous human decisions are modeled using statistical tools such as softmax and discrete, and Parzen and Gaussian sum distributions. Statistical formalism is maintained in order to enable probabilistic analysis and prediction. The theory has been applied to human decision data using RoboFlag, a multiple robot simulator of capture the flag; data was collected in a series of three tests, jointly developed and implemented with AFRL/HECP. Program contributions summarized here include: 1) probabilistic modeling of both discrete and continuous human decisions using probabilistic graphs; 2) Coupled human-vehicle probabilistic models; 3) empirical studies of human-vehicle interaction with a focus on adaptive tasking; 4) theory and experiments with formal human sensor networks; 5) cooperative geolocation using uninhabited aerial vehicles.					
15. SUBJECT TERMS mixed initiative control, human-vehicle interaction, human decision modeling, cooperative geolocation, probabilistic graph models					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Mark E. Campbell
U	U	U	UU		19b. TELEPHONE NUMBER (Include area code) (607)255-4268

20090914177

Modeling for Semi-Autonomous Control of Cooperative Vehicles

Prof. Mark E. Campbell

Cornell University

Final Report for AFOSR Contract #: FA9550-05-1-0118

Abstract

This report details work completed in an AFOSR program focused on semi-autonomous control of cooperative vehicles. Coupled operator-multiple vehicle systems are modeled in a unified framework using probabilistic graphs to yield a methodology for analyzing semi-autonomous systems. The framework uses conditional probabilistic dependencies between elements, leading to a Bayesian network with probabilistic evaluation capability. Both discrete and continuous human decisions are modeled using statistical tools such as softmax and discrete, and Parzen and Gaussian sum distributions. Statistical formalism is maintained in order to enable probabilistic analysis and prediction. The theory has been applied to human decision data using RoboFlag, a multiple robot simulator of capture the flag; data was collected in a series of three tests, jointly developed and implemented with AFRL/HECP. Program contributions summarized here include: 1) probabilistic modeling of both discrete and continuous human decisions using probabilistic graphs; 2) Coupled human-vehicle probabilistic models; 3) empirical studies of human-vehicle systems with a focus on adaptive tasking; 4) theory and experiments with formal human sensor networks; 5) cooperative geolocation using uninhabited aerial vehicles.

Report Outline

The following five areas of contributions, as developed within this program, are summarized in this report:

- 1) probabilistic modeling of both discrete and continuous human decisions using probabilistic graphs;
- 2) Coupled human-vehicle probabilistic models;
- 3) Empirical studies of human-vehicle interaction with a focus on adaptive tasking;
- 4) Theory and experiments with formal human sensor networks;
- 5) Cooperative geolocation using uninhabited aerial vehicles.

Each section also gives a summary of publications that were partially or fully funded by this program. A subset of these publications are appended to this report.

The final two sections of this report summarize the technology transfer and educational contributions from this program.

Table of Contents

Contribution #1: Human Decision Modeling Theory	3
Contribution #2: Coupled Human-Vehicle Probabilistic Models.....	10
Contribution #3: Empirical Studies of Adaptive Tasking in Human-Vehicle Interaction with AFRL/HECP	11
Contribution #4: Human Sensor Networks.....	15
Contribution #5: Cooperative Estimation and Control	18
Technology Transfer	20
Education.....	20
Appendix: Partial Publications from the Program	21

Contribution #1: Human Decision Modeling Theory

This contribution is to bring formal, probabilistic models to bear on the problem of human decision modeling. Incorporation of the human element is critically important for future defense missions, yet research in this area has had a relatively soft theoretical foundation.

This is in contrast to vehicle technologies (control, tasking, tracking, etc.) which are typically quite formal. If formal methods can be developed for human decision modeling, a large amount of other work could be leveraged for human-automation systems (e.g. networked control, task allocation, etc.)

Our approach has been to model human decisions as a probabilistic graph. Probabilistic graph models, also known as Bayesian Network (BN) models and Dynamic Bayesian Network (DBN) models (with time dependency), will be developed for human decisions, tasking, and sensory inputs. These models nicely capture variations across subjects and environmental conditions, intuitively describe decision dependencies, and can model joint human-vehicle systems – all in a formal probabilistic framework. The framework is ideal for modeling human-human collaboration as well as hierarchical decisions.

To more easily understand the capabilities of these models, consider Figure 1, which is an integrated model of an operator commanding a UAV to track ground targets using a camera. The probabilistic model is represented by a directed graph, with each arrow indicating conditional probabilistic dependency. For example, the UAV position sensor (GPS) is nicely modeled as a normal Gaussian distribution, or $p(Y_{uav}|X_{uav}) = N(X_{uav}, \Sigma_v)$. Note that the model in Figure 1 is a BN with no time dependence; a Dynamic Belief Network (DBN) model is developed by repeating the BN model over a series of time slices, with an appropriate time based model. By maintaining connections between nodes, the model can be used for formal probabilistic algorithm development in estimation, prediction and optimization.

Figure 1 shows two models which are well known in autonomous vehicles: Pose and Target Tracking. The Uninhabited Aerial Vehicle (UAV) model includes the mode of operation (U), position/attitude state (X_{uav}), and sensor measurements (Y_{uav}). Estimation of the vehicle state can be accomplished using the well known Kalman Filter, which can be written in terms of the graph as $p(X_{uav}|U, Y_{uav})$ ("probability of the UAV state, given the mode of operation and measurements"). The Target Tracking model includes the target type (V), position state (X_{tar}), and camera measurement (Y_{cam}), which depends on the location of the target and location/attitude of the UAV. Joint estimation of the target type and state is a classic target tracking problem, which can be solved with a multiple model Kalman Filter, $p(X_{tar}, V | X_{uav}, Y_{cam})$.

The third model, which is the focus of this work, is the Operator Decision Model. In this example, the decision (D) is defined as the discrete tasking of the UAV by the operator (i.e. selection of the UAV mode of operation U at the next time step). This decision is naturally a probabilistic function of the system states: $p(D|X_{uav}, V, X_{tar}, U)$, or the "probability of a Decision choice is dependent on the UAV mode and location; target type and location."

Key to bringing control theoretic algorithms to the proposed human-vehicle network is the ability to model human operator decisions. Our work has developed a variety of methods for identifying (learning) probabilistic graph models for both discrete and continuous human decisions.

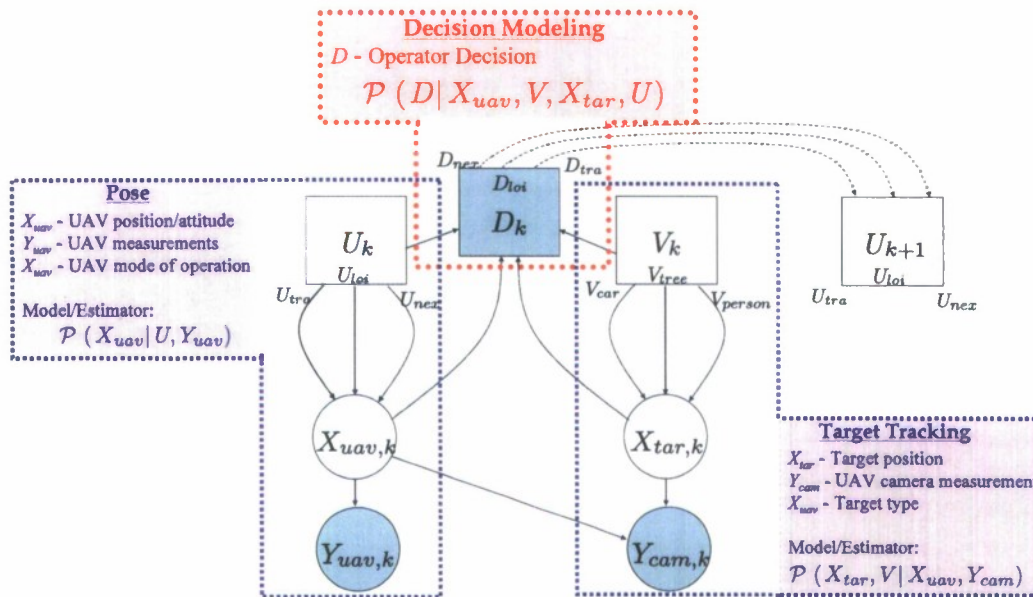


Figure 1: Probabilistic graph model of a single operator tasking a single UAV to track (and estimate states) of a target. Circles denote continuous vector valued random variables, squares discrete random variables, and shaded blocks indicated that direct measurements are available.

Our work has produced a general approach to modeling both discrete and continuous human decisions, given a set of training data[1]-[7]. The operator decisions are modeled in this framework by defining and identifying $p(D|X)$, which is the likelihood of the decision data (D) given the system state variables (X). In the discrete case, as defined in Ref. [3], a “decision model block,” which includes internal nodes of logistic or softmax thresholding distributions. This notion is represented in Figure 2(left). Ref. [3] also outlines a solution procedure for this joint density in order to find a common, conditional density $p(D|X)$, for any block structure. The simplifies to:

$$\mathcal{P}(D|X) = \sum_{D^{(1)}:D^{(s)}} \prod_{l=1}^s \mathcal{P}(D^{(l)}|X) \cdot \mathcal{P}(D|D^{(1)}:D^{(s)})$$

where the term on the right are the internal softmax distributions, and the term on the left is a discrete selection matrix. This important conclusion enables the optimization of internal distribution parameters as well as the “structure” (graph and connections) under one common cost function. Figure 2(right) shows how the optimization notionally works, where the discrete selection matrix is fixed, and the distribution parameters of the softmax variables are optimized. Once the likelihood begins to level off with increased numbers of softmax variables, additional “clusters” of graphs are added in order to improve the identified fit.




Figure 2: Bayesian Network determination. Left: Graph of a general decision block $P(D|X)$ showing internal nodes of the graph. Right: Optimization of the Likelihood as a function of internal softmax variables and discrete selections.

A principle feature of this approach is that it produces linear boundaries between decision clusters in state-space. The number of these linear boundaries is given by the number of hidden logistic nodes used inside the BN. In addition, since the BN node distribution parameters are obtained using maximum likelihood estimation (MLE), the likelihood function can be used to gauge model accuracy and over-fitting as more nodes are added to the BN; this also allows the fitting process to be automated off-line. Ref. [3] showed that with infinite data, the optimized parameters converge to the true distribution; with finite data, the parameters are bounded by confidence bounds in the form of a Chi-square distribution. The logistic BN classifier is a binary classifier naturally suited for discriminating between convex, unimodal clusters of decision data; for convex and unimodal multinomial decision data with N decision classes, N separate BN classifiers can be identified in parallel. Non-convex or multi-modal decision clusters, however, are more difficult to classify in this approach due to the nature of the logit probability function.

The methodology in Ref. [3] was then evaluated using data from a set of RoboFlag experiments. Operators controlled three vehicles: two search vehicles, each with the ability to locate entities; and one ID vehicle, with the ability to identify the entity's type. During the mission, three types of entities could be encountered: a stationary flag; a stationary red robot which can tag if blue vehicles come too close; and a red chase vehicle, designed to chase any blue vehicle within a particular range. Probability of localization (uncertainty circles) and identity (color or circle) were improved when vehicles moved closer to the targets. When users are confident of the final type (classification), they formally selected the target type using a GUI.

Results for 16 operators summarizing a simple case of tasking the ID vehicle tasking are shown in Figure 3(left). Also, shown are dotted lines which represent the decision boundary, as defined by the probabilistic BN. Figure 3(right) shows a typical probability plot for the BN, $p(D=Loiter|X)$. Results show that users typically tasked the ID vehicle to move and identify the target when the ID probability ~ 0.5 (i.e. unknown). The loiter mode was commonly used at the start of the game (ID probability ~ 0.5 , location uncertainty large), or in the middle of the game when moving to wait for the next target (ID probability ~ 1 , location uncertainty small).

One of the more interesting cases is at the center of Figure 3(left), where users tasked the ID vehicle to move in and ID the target *before* the location uncertainty had decreased (i.e. the users were still very unsure where the target was). During interviews and in the user videos, it was realized that many people performed a “strategic” play, when they sent both the localizing and ID vehicles in towards the target at the same time. The probabilistic models in Ref. [3] cannot, unfortunately, recognize the coupling of these moves over time. This is an important area of future research.




Figure 3: Results from AFRL led RoboFlag experiments for tasking the ID vehicle; operator decisions were classified as either tasking the ID vehicle to loiter, or move in and ID the target. Right: Likelihood graph after the BN has been identified.

These modeling concepts were extended in theory for both the continuous and discrete model cases. In the discrete case, the primary extension was to develop a more general modeling approach that considered multi-modal human decisions [6]-[7]. Conventional softmax regression can be extended to multi-modal decision data by subdividing the data prior to model identification. Consider the case of n_D discrete decisions in D , and a set of system state variables X . Assuming that the data for the i th operator decision D_i can be clustered into d_i exclusive subsets, then d distinct softmax conditional probabilities can be estimated similar to the approach in Ref. [3]. The operator decision likelihood is then simply the sum of the corresponding d_i softmax probabilities.

$$\mathcal{P}(D = i | \mathcal{X}) = \sum_{j=1}^d \mathcal{P}(D = i | S = j) \cdot \mathcal{P}(S = j | \mathcal{X}) = \sum_{\ell=1}^{d_i} \mathcal{P}(D = i | S = \ell) \cdot \mathcal{P}(S = \ell | \mathcal{X})$$

The resulting BN is shown in Figure 4, which shows extensions to multi-modal discrete decision data. Figure 4(center) shows data for $n_D=3$ decisions; this data can be further divided into d_i subdivisions, and Figure 4(right) illustrates the estimated BN. Note that the likelihood $p(D=1|X)$ no longer ensures a convex optimization problem for estimating the softmax weights, as it was in the original procedure. However, the Hessian matrix for the weights is never positive definite, so that the likelihood can only have local maxima or saddle points, thus most likely yielding good fits to the data.




Figure 4: Multi-modal modeling: Left: The BN model. Center: Two-dimensional data set with $n_D=3$ discrete decisions. Right: The conditional probability for one decision,

For more complex decision models, which may include complex decision boundaries (in the case of discrete decisions) or continuous decisions, non-parametric modeling procedures were developed. In this case, class-conditional probability density functions (pdfs) $f(X|D)$ for the state-variables are estimated, and coupled with prior probabilities $p(D)$, then an indirect estimate of $p(D|X)$ can be developed using Bayes' rule:

$$P(D = i|X) = \frac{f(X|D = i) \cdot P(D = i)}{\sum_{j=1}^{n_D} f(X|D = j) \cdot P(D = j)}$$

The decision class priors $p(D)$ may be estimated from the training data using frequency counts, while class-conditional probability densities $f(X|D)$ can be estimated via parametric, semi-parametric, or non-parametric methods. Parzen's kernel method is widely used for nonparametric density estimation in statistical classification because the estimates are easy to train and can offer robust performance without making any assumptions about the underlying class-conditional densities. For a given decision class i with training set x_i with n_i points, the Parzen density estimate is given by a sum of n_i kernel pdfs centered about each sample point. Gaussian kernels are common, leading to a sum of Gaussians mixture model.

The nonparametric approach leads to the simple BN representation shown in Figure 5. Since Parzen density estimates asymptotically approximate any probability density function, they are ideal for generating non-linear decision surfaces. Figure 5(center) shows a sample training set for $n_D=3$ decisions that are non-linearly correlated, and Figure 5(right) shows the post-processed likelihood function, $p(D=2|X)$ for the data.

In addition to the general theoretical approach, three significant theoretical guarantees were made in this work:

- Given a set of labeled, discrete human decision data, we have developed an approach that guarantees convergence to the true distribution parameters in the limit of unlimited data.
- Given a set of labeled, discrete human decision data in non-convex clusters, no matter how large, we have developed an approach that guarantees convergence to the true distribution parameters in the limit of unlimited data.
- Given a set of labeled, continuous human decision data, we have developed a mixed Gaussian representation with a new factorization that enables ease in identifying the distribution parameters; in the limit of unlimited Gaussians and data, the approach is guaranteed to approach the true distribution.

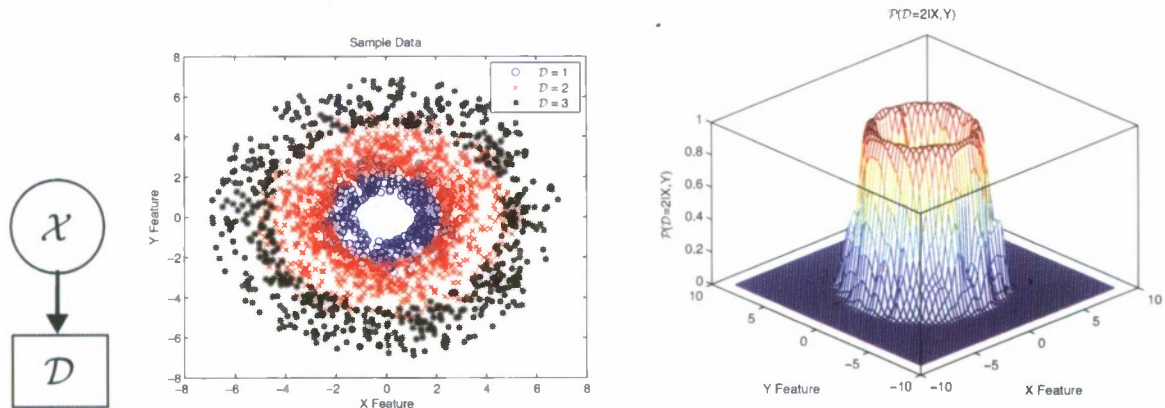


Figure 5: Nonparametric modeling: Left: The BN model. Center: Two-dimensional data set with $n_D=3$ discrete decisions exhibiting non-linear class correlations. Right: Post-processed conditional probability for one decision.

Extensive evaluation of the modeling procedures has taken place. Because this is a system identification modeling procedure (i.e. developing models from data), a data holdout approach was taken. In addition, our approaches were generally benchmarked against state of the art research, and with other datasets that are accepted in the field. The multi-modal decision modeling approach, termed the multimodal softmax (MMS) model, exploits the concept of discriminative subclass modeling to estimate parameters for complex class posterior distributions without relabeling training data, in contrast to existing 'hard subclass' discriminative modeling strategies. The MMS model can also be shown to provide a discriminatively learnable generalization of both the softmax model and the widely used class of Gaussian mixture model based posterior models. Classification experiments with benchmark and application data compare the training cost and accuracy of the MMS model with those of hard subclass models (HSS) and other well known probabilistic models, such as kernel methods and hierarchical mixture of experts (HME). The MMS model is shown to achieve a high level of discriminative modeling accuracy with relatively low training cost:

TABLE 1
Benchmark Classification Error Statistics (mean % error, standard deviation).

Data Set	Softmax	MMS	HSS	Best ME
Glass	38.10 ± 2.14	24.10 ± 3.33	34.76 ± 2.67	32.93 ± 2.47 (ME2)
WBCO	3.34 ± 0.19	2.12 ± 0.19	3.29 ± 0.27	3.85 ± 0.28 (ME3)
Banana (linear)	47.09 ± 4.52	12.32 ± 2.89	13.17 ± 1.37	12.76 ± 0.85 (ME4)
Banana (quadratic)	35.94 ± 3.79	11.37 ± 0.75	15.20 ± 3.82	11.80 ± 0.83 (ME3)
Thyroid	13.01 ± 5.00	7.12 ± 4.13	7.28 ± 4.25	9.11 ± 4.37 (ME6)
Image	48.16 ± 9.28	6.51 ± 2.26	45.03 ± 8.88	48.16 ± 9.27 (ME3)
Twonorm (linear)	4.20 ± 0.69	4.59 ± 0.67	5.68 ± 1.01	3.90 ± 0.43 (ME3)
Twonorm (RBF)	4.65 ± 1.91	3.36 ± 0.58	4.57 ± 1.93	3.77 ± 0.87 (ME6)

TABLE 2
Benchmark Training Costs (mean training time (CPU secs), parameter count).

Data Set	Softmax	MMS	HSS	Best ME
Glass	0.90, 60	1.55, 130	2.78, 130	23.64, 140
WBCO	1.24, 20	0.15, 30	1.60, 30	68.98, 90
Banana (linear)	0.60, 6	2.37, 24	1.15, 24	40.35, 36
Banana (quadratic)	0.67, 12	0.71, 24	1.15, 24	38.62, 54
Thyroid	0.40, 12	0.28, 18	0.45, 18	18.76, 108
Image	4.38, 42	27.55, 84	6.52, 84	151.63, 189
Twonorm (linear)	1.46, 21	0.81, 84	2.10, 84	50.83, 189
Twonorm (RBF)	1.00, 6	2.21, 12	1.20, 12	48.73, 54

The decision modeling approach was also evaluated using human decision data from the RoboFlag experiments, where humans tasked robots in a search and ID mission. The classification problem attempted to find the probability of a decision D, given a 12 dimensional input feature vector X. All models were trained and validated via 30 holdout trials. Table 3 shows the classification errors for each model on each of the three example data sets; statistically significant comparison results are shown in bold. Table 4 shows the training time. In all cases, the MMS model appears to strike a very good balance between training time and accuracy.

TABLE 3
RoboFlag Classification Error Statistics (mean % error, standard deviation).

Case	Softmax	SVM	MMS	GMM-A	GMM-B	HSS-A	HSS-B	Best ME
1	19.09 ± 0.63	15.71 ± 0.48	17.11 ± 0.74	17.91 ± 0.69	16.83 ± 0.47	18.35 ± 0.95	19.17 ± 1.03	16.27 ± 0.43 (ME2)
2	23.23 ± 0.68	18.53 ± 0.83	18.01 ± 0.85	21.40 ± 0.77	20.65 ± 0.77	19.36 ± 0.77	19.21 ± 0.74	18.47 ± 0.68 (ME3)
3	61.93 ± 0.28	40.62 ± 1.11	40.25 ± 1.06	43.12 ± 1.76	39.69 ± 1.51	41.93 ± 1.13	43.22 ± 1.11	41.42 ± 0.91 (ME4)

TABLE 4
RoboFlag Training Costs (mean training time (CPU secs), parameter count).

Case	Softmax	SVM	MMS	GMM-A	GMM-B	HSS-A	HSS-B	Best ME
1	2.24, 39	173.06, 1222	16.21, 143	0.56, 13167	0.90, 21546	11.61, 143	34.51, 234	71.93, 104
2	1.89, 39	135.52, 1404	5.30, 104	0.37, 9576	0.37, 10773	6.06, 104	7.86, 117	96.47, 156
3	4.63, 78	40.53, 1428	9.42, 143	0.27, 13167	0.73, 21546	18.45, 143	48.96, 234	219.48, 208

Publications related to this contribution, and funded in part by this program, include the following. Note that the publications with a '*' are attached to this report in an appendix.

- [1] J. Veverka, M. Campbell, "Towards an Operator Decision Model for ISR Type Missions," *AIAA Guidance, Navigation and Control Conference*, Aug 2005.
- [2] J. Veverka, M. Campbell, "Operator Decision Modeling for ISR Type Missions," *IEEE Conference on Systems, Man, and Cybernetics*, Oct 2005.
- [3] *M. Campbell, S. Sukkarieh, A. Goktogan, "Operator Decision Modeling in Cooperative UAV Systems," *2006 AIAA Guidance, Navigation and Control Conference*.
- [4] M. Campbell, N. Ahmed, S. Sukkarieh, "Operator Decision Modeling in Cooperative UAV Systems," submitted to the *AIAA Journal of Guidance, Control and Dynamics*.
- [5] D. Shah, M. Campbell, "State-Dependent Probabilistic Model Reduction for Evaluation of Human-Robotic Autonomous Systems", *2007 IEEE Systems, Man and Cybernetics Conference*.
- [6] N. Ahmed, M. Campbell, "Multimodal Operator Decision Models," *2008 American Control Conference*, Seattle WA.
- [7] *N. Ahmed, M. Campbell, "Discriminative Subclass Modeling without Relabeling," submitted to the *IEEE Transactions on Systems, Man and Cybernetics*.

Contribution #2: Coupled Human-Vehicle Probabilistic Models

This contribution was to explore how far a human-vehicle system could be jointly modeled using the probabilistic graph theory. While this contribution was largely systems like in nature, it provides a solid understanding as to where the advantages and limitations in modeling these systems are. Initial work has been completed in Refs.[8],[9], while further work, including an evaluation using data holdout, is on-going.

The RoboFlag experiments provide large sets of data which can be used to explore coupled operator-vehicle modeling. One example is to decompose the system into hierarchy of control/decision making, as shown in Figure 6(left). Four levels are assumed: Vehicle(s) Selection and coordination; Play Selection for strategies; User Clicks for tactical locations; Vehicle Motion based on its actuation. The vehicle motion model is deterministic. All models based on human decisions are developed using the techniques from our research, based on the data.

As an example of developing a probabilistic model from operator data and using it for prediction, consider the case of an operator tactically tasking the search vehicle (i.e. the operator selecting a point on the field for the search vehicle to move). This is the shaded block in the BN in Figure 6(left). Figure 6(center) illustrates the entire set of tactical search data obtained from the RoboFlag I data (vehicle start points, assignments, and trajectories). A probabilistic model of this decision can be developed using the nonparametric Parzen density estimation, as described earlier. Figure 6(right) shows how this is constructed and used. The vehicle is at the blue star. From the RoboFlag data, previous cases of users tasking the search vehicle from a "similar location" (within a particular probability level of the nominal location; these are denoted by {red stars}. The corresponding vehicle assignment points, denoted by {black dots}, are each assigned a Gaussian kernel at the location of the {black dots}. The vehicle assignment pdf is then formed by summing these Gaussians, weighted by the location of the starting point (denoted by the height of the red points). Figure 6(right) shows the density contours of this pdf, which could be used to find the corresponding likelihood, $p(U_{\text{tact}}|U_{\text{strat}}, X)$ in Figure 6(left).




Figure 6: Coupled operator-vehicle model based on RoboFlag. (a) Hierarchical graph model; (b) Data for user tasking search vehicles (start: {red stars}, way point assignment: {black dots}, and path: line); (c) Contour of the probabilistic tactical decision density model.

Publications related to this contribution, and funded in part by this program, include the following. Note that the publications with a '*' are attached to this report in an appendix.

- [8] *F. Bourgault, N. Ahmed, D. Schrader, M. Campbell, "Probabilistic Operator-Multiple Robot Modeling Using Bayesian Network Representation," *2007 AIAA Guidance, Navigation and Control Conference*.
- [9] M. Campbell, F. Bourgault, S. Galster, D. Schneider, "Towards Probabilistic Operator-Multiple Robot Decision Models," *2007 International Conference on Robotics and Automation*.

Contribution #3: Empirical Studies of Adaptive Tasking in Human-Vehicle Interaction with AFRL/HECP

This contribution was to empirically study, with collaborators from AFRL/HECP, the concept of adaptive tasking in human-vehicle systems; the data from these tests could then be used to validate the modeling approaches (described in Contributions 1 and 2).

The approach was to utilize a testbed developed at Cornell, called RoboFlag. RoboFlag is a human controlled simulator of the game 'Capture the Flag.' RoboFlag is ideally suited to study semi-autonomous systems because operators can be easily integrated into the loop, basic mission types such as ISR can be studied, and system parameters and uncertainties can be manipulated quickly (such as communications constraints, vehicle types, and sensor types). The simulator has been used extensively by AFRL and the Co-PI's group as a basic research tool.

A modified game was developed for this project, in conjunction with Drs. Scott Galster and Benjamin Knott from AFRL/HE. The goal was to locate and identify two targets, which could be flags or red agents. Operators controlled three vehicles: two fast search vehicles, each with the ability to locate entities; and one slower ID vehicle, with the ability to identify the entity's type. During the mission, three types of entities could be encountered: a stationary flag; a stationary red robot which can tag if blue vehicles come too close; and a red chase vehicle, designed to chase any blue vehicle within a particular range. Probability of localization (uncertainty circles) and identity (color or circle) were improved when vehicles moved closer to the targets. When users are confident of the final entity type (classification), they formally selected the target type using a GUI input.

Three sets of games were conducted: manual (way point control) in 2006, automated (plays developed from the manual data) in 2007, and adaptive in 2008. The differences in the games focused on the interactions and feedback loops. In the first set, the operator primarily controlled the vehicles *manually*, with point and click waypoints. In the second set, there was a suite of *automated* plays derived from our understanding of the manual game. In the third, *feedback* loops were added during and after the games, based on the primary metrics of the game (see below). For each study, 16 subjects from the AFRL subject pool were used, with 24 trials each; a 4x4 matrix of target location and type was used. All vehicle telemetry data was saved, and users were video recorded. GUIs for these three games are shown in Figure 7.




Figure 7: GUIs for the three sets of RoboFlag games. From left to right: Manual, Autonomous, and Adaptive.

With all three sets of games completed, recent work has focused on statistically analyzing the results of the games, and across the games. For the analysis across each game, six metrics were statistically analyzed to evaluate the users' mission performance:

Mission time: total time spent completing each trial; typically correlated to user efficiency. Shorter mission times indicate that the human user was able to allocate resources efficiently and implement appropriate actions.

Idle time: percent of time robots were without instruction; typically caused by high operator load. Robots were considered idle when they were without instruction (standing still), except in cases where they were collecting location or ID information from an in-range target. In general, an idle robot indicates an underutilization of resources because it is not actively contributing to the mission objectives. This typically occurs when the operator is distracted (for example, a robot is being chased by the CHS) and is unable to allocate attention to all robots equally.

Tag events: number of times a robot was 'tagged' by an enemy agent. Robots are temporarily unusable when tagged, so these are obviously undesirable events. Not only do tag events decrease the resources available to the operator, but they may necessarily increase the overall mission time if targets can not be located/identified until the tagged robot(s) become active again. This is especially the case when the tagged robot is the sole IDV.

Human input: frequency of instructions given by human user; typically a measure of operator workload/attention. The number, frequency, and type of instructions given by the operator are measures of mental workload and attention allocation.

Target ID probability: amount of identity information collected when target identified.

Target location radius: amount of location information collected when target identified.

To test the effects of the experimental manipulations on the above performance metrics, a 3 (RF game number) x 2 (flag location -- near/far) analysis of variance (ANOVA) was performed, the results of which are presented in the table below:

Table 1: Results of 3 x 2 ANOVA, where significant interaction effects are reported.

Mission time	$F(2, 547) = 5.80,$	$p < 0.01$
% Idle Time	$F(2, 547) = 27.48,$	$p < 0.001$
Num. tag events	$F(2, 547) = 7.76,$	$p < 0.001$
Num. user inputs	$F(2, 547) = 5.70,$	$p < 0.01$
ID probability	$F(2, 547) = 18.31,$	$p < 0.001$
Location radius	$F(2, 547) = 4.42,$	$p < 0.05$

This table shows that there were significant interaction effects across all six metrics. These interactions more fully using post hoc paired t-tests. While the full results are summarized in Ref. [12], the results for 'Mission Time' are summarized here.

The results of a 3 (game number) x 2 (flag location) ANOVA indicate that there was a significant interaction between these factors on average mission time, $F(\{2,547\}) = 5.80, p < 0.01$. A post hoc paired t-test was performed to determine the source of variance, the results of which indicate that there was a significant difference in mission time between all three games, as shown in this table

Table 2: Results of post hoc paired t-test for mission time.

I vs. II :	$t(367.14) = 2.042,$	$p < 0.05$
I vs. III :	$t(246.10) = 7.201,$	$p < 0.001$
II vs. III :	$t(305.13) = 6.587,$	$p < 0.001$

It can be seen in Figure 8 that users took longest on average to complete set I trials, where robots could only be controlled using waypoints. The addition of automated plays (especially Evade/Avoid) in set II decreased overall game time by an average of 11.40 seconds, almost 10%. In set III, users could use both automated plays and waypoint control, and performance feedback (especially the idle-alert feature) encouraged greater efficiency. As a result, the average set III mission time decreased another 22.44 seconds -- approximately 21% from set II.

A post hoc paired t-test also showed a significant difference ($t(\{202\}) = 3.036, p < 0.01$) between near/far target configurations for game I, but not for games II or III. This indicates that increasing automation had the desirable effect of increasing uniformity in human efficiency both between mission configurations and across users.

Interestingly, mission time in *near* cases (targets located close to one another near the Home Base) was not significantly improved with the addition of automated plays, suggesting that these trials were 'easy' enough that waypoint-only control did not hinder the users' ability to perform the mission efficiently. However, for *far* cases (targets located far from Home Base and from one another) average mission time improved by about 25.6 seconds (19.5%) with the addition of automated plays ($t(\{168.30\}) = 3.323, p < 0.01$). Users were able complete these 'harder' missions in less time with assistance from automatic Evade and Avoid decreasing tag occurrences, and by using strategic plays like Decoy to give more complicated commands without using distracting and time-consuming waypoint control. Average mission time decreased an additional 23.7 seconds (22.3%) when performance feedback was provided ($t(\{161.48\}) = 4.879, p < 0.001$), as users were kept aware of idling robots and end-of-game scoring encouraged lower mission times.




Figure 8: Mission time for near/far trials of all three RoboFlag experiments, units are seconds. The boxes have lines at the lower quartile, median, and upper quartile values. The whiskers extend from the box out to the most extreme data value within 1.5xIQR (interquartile range of the sample). The circle-cross symbols are located at the sample means.

Publications related to this contribution, and funded in part by this program, include the following. Note that the publications with a '*' are attached to this report in an appendix.

- [10] D. Schneider, M. Campbell, S. Galster, A. Klochko, J. Veverka, "The RoboFlag Test System for Decentralized Autonomous and Semi-Autonomous Cooperative Multi-Agent Systems Research," submitted to the *IEEE Transactions on Control Systems Technology*.
- [11] D. Shah, S. Galster, M. Campbell, F. Bourgault, N. Ahmed, B. Knott, "A Study of Human-Robotic Teams with Various Levels of Autonomy," *2009 AIAA Infotech conference*.
- [12] *D. Shah, S. Galster, M. Campbell, N. Ahmed, B. Knott, "On Operator Control and Human-Robotic Team Performance," submitted to the *IEEE Transactions on Systems, Man and Cybernetics*.

Contribution #4: Human Sensor Networks

This contribution was to develop formal theory and perform experiments of sensor networks with human elements. Typical sensor networks are large, yet the information is very simple; many times, such as in consensus methods, probabilities are not even taken into account. The challenge with humans is that the information will be complex; the probabilistic graph theory, which fits well here by developing decision likelihoods, creates models that are intuitive, but complex.

The general approach to sensor networks and fusion is shown in this equation:

$$\underbrace{p(\mathbf{x}_k | \mathbf{I}_k)}_{\text{Fused Search PDF}} = \underbrace{p(\mathbf{x}_k | \mathbf{I}_{k-1})}_{\text{Predicted Search PDF}} \prod_{i=1}^{N_s} \underbrace{p(\mathbf{z}_k^i | \mathbf{x}_k)}_{\text{Measurement Likelihoods}}$$

Here, the priori pdf is fused with measurements (and likelihoods) from different sensors.

Two elements of the theory were explored in this case. The first was modeling humans as sensors. In typical sensors, models are termed probabilistic likelihoods, which are usually simple, such as Gaussian distributions. But, they are not a function of time, which allows the BN probabilistic graph model work developed here to be used. The second element of the theory is the information sharing. Our approach was to use a grid based approach, so that the general likelihood models from humans could always be considered (at the expense of computation). Channel filtering concepts were used, which were first developed using Kalman Filter based theory. Here, each communication channel is monitored for information sent back and forth. Only new information is sent at each time step. For tree-like topologies, even in the presence of communication uncertainties such as drops, the solution is guaranteed to converge to the centralized solution. Details of the theory can be found in Refs. [13].[14].

The applicability and limitations of the operator decision modeling theory for a realistic scenario was tested using humans in a search experiment. The mission was to task a group of humans to search for a white golf ball on Cornell's Ithaca campus, in the winter. Five human operators were equipped with handheld computers which were linked wirelessly in an ad hoc network. This network would drop packets, and not communicate when out of range; therefore, the fusion process had to include these assumptions in order to maintain theoretical consistency in the estimator. All computers had (uncertain) GPS and compass measurements; so, likelihoods for these measurements were also available.

A BN model of a human 'search' measurement likelihood (target detection) was developed, as shown in Figure 9(left), which shows dependency on range and bearing. Figure 9(center) shows the likelihood of target detection using human vision. This model was developed using actual measurements of instances of detection/no detection for several human subjects. The model shows a forward-facing human, where maximum target visibility is straight ahead. As seen in Figure 9 (left), detection drops off and is approximately zero when the target is located >3m from the human sensor, mostly due to snow, shadows, or other visual obstructions. Target detection drops off on the sides as peripheral vision makes it difficult for the human to see possible targets.




Figure 9: Likelihood of Detection models for human vision (left,center) and robot (right). The vertical axis (y) is along the human's center of gaze. The horizontal axis is at the edge of the human's peripheral vision, where there is no target visibility. Red contour lines indicate high likelihood of detection. All units are in meters.

Figure 10 shows the experiment hardware (left) and results for one node (right). All fusion was accomplished using a grid based PDF, which was used to present the user with a probability of target location. Because grid based methods were used, constraints such as buildings could easily be incorporated. Figure 10(right) shows the results from one human node, which overlays a satellite image of campus with the real time fused PDF of the target location; also shown in the trajectory of one human's path. The human moves towards locations of high target location probability; as the human moves, its measurement (GPS, compass/heading, target yes/no) is fused into the network and shared appropriately. The modeling and fusion process worked well, even for a system with up to five nodes [13].[14].

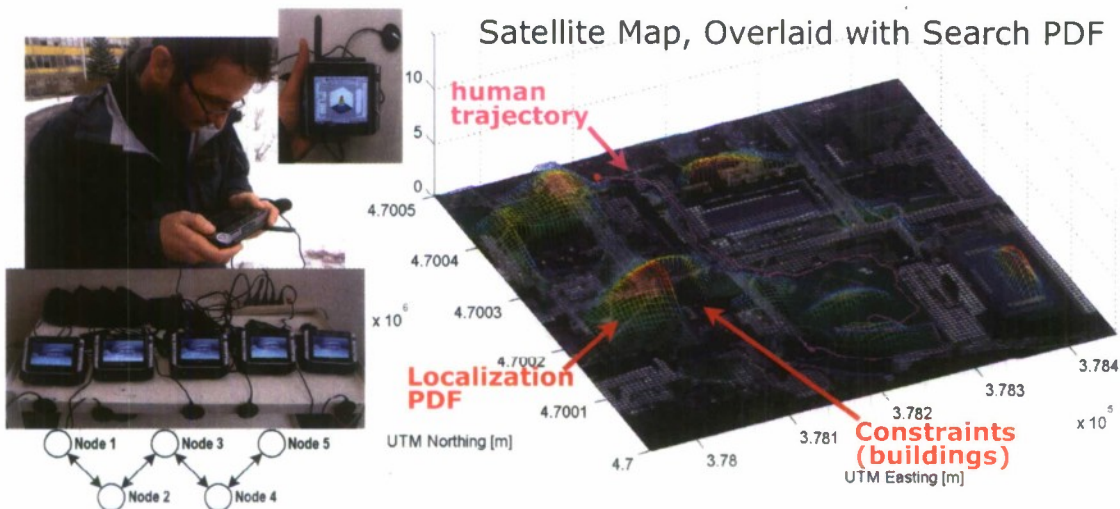


Figure 10: Recent experiment with a network of humans in a probabilistic search experiment. Left: Hardware, including five handheld PC's with local network capabilities. Right: Overlay of satellite Imagery with a probability density of "probable" locations of the target.

Publications related to this contribution, and funded in part by this program, include the following. Note that the publications with a '*' are attached to this report in an appendix.

- [13] F. Bourgault, A. Chokshi, M. Campbell, "Human-Computer Augmented Nodes for Scalable Mobile Sensor Networks," *2008 IEEE SMC International Conference on Distributed Human-Machine Systems*, Athens Greece.
- [14] *F. Bourgault, A. Chokshi, J. Wang, D. Shah, F. Cedano and M. Campbell, "Scalable Bayesian Human-Robot Cooperation in Mobile Sensor Networks," *2008 IEEE International Conference on Intelligent Robots and Systems*.
- [15] F. Bourgault, A. Chokshi, J. Wang, D. Shah, F. Cedano and M. Campbell, "Scalable Bayesian Human-Robot Cooperation in Mobile Sensor Networks," to be submitted to the *International Journal of Robotics Research*.

Contribution #5: Cooperative Estimation and Control

While this project primarily focused on human decision modeling, the application of interest was Intelligence, Surveillance and Reconnaissance missions (ISR). Thus, some work in autonomous UAVs (single and cooperative) for missions such as geolocation, was performed [16]-[30]. Contributions here range from verification of on-line estimation and control algorithms on UAVs; path planning theory which maximizes information collection across multiple vehicles; and cooperative geolocation (theory and experiments) using UAVs with gimbaling camera payloads. The latter was the primary focus, with research in information fusion work, cooperative bias estimation, and cooperative planning for tracking.

Publications related to this contribution, and funded in part by this program, include the following. Note that the publications with a '*' are attached to this report in an appendix.

- [16] D. Schneider, M. Campbell, "Real Time Optimal Task Allocation in Highly Dynamic Environments," *ASME International Mechanical Engineering Congress and Exposition*, Nov 2005.
- [17] M. Campbell, J. W. Lee, E. Scholte, D. Rathbun, "Flight Results for On-line Estimation, Planning and Control using the SeaScan UAV," *AIAA Journal of Guidance, Dynamics and Controls*, Vol 30, No 6, 2007, pp. 1597-1609.
- [18] E. Scholte, M. Campbell, "Robust Nonlinear Model Predictive Control with Partial State Information," *IEEE Transactions on Control System Technology*, Volume 16, No. 3, 2008.
- [19] *M. Campbell, W. Whitacre, "Cooperative Tracking using Vision Measurements on SeaScan UAV's," *IEEE Transactions on Control System Technology*, Vol. 15, No. 4, July 2007, pp. 613-626.
- [20] W. Whitacre, M. Campbell, "Cooperative Estimation in Networks of UAV's with Delayed Data," *2007 AIAA Guidance, Navigation and Control conference*.
- [21] J. Ousingawat, M. Campbell, "Planning for Cooperative Multi-vehicle Reconnaissance," *AIAA Journal of Aerospace Computing, Information, and Communication*, Vol. 4, No. 1, Mar-Apr 2007.
- [22] J. Ousingawat, M. Campbell, "Optimal Planning for Cooperative Reconnaissance Using Multiple Vehicles," *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 1, Jan-Feb 2007, pp. 122-132.
- [23] D. Stevenson, M. Wheeler, M. Campbell and W. Whitacre, R. Rysdyk and R. Wise, "Experiments in Cooperative Tracking of Moving Targets by a Team of Autonomous UAV's", *2007 AIAA Guidance, Navigation and Control Conference*, AIAA-2007-6756
- [24] M. Wheeler, R. Wise, R. Rysdyk, W. Whitacre, M. Campbell, "Autonomous Cooperative Geo- Location and Tracking of Moving Targets," *2007 Infotech@Aerospace Conference*, AIAA-2007-2852.
- [25] W. Whitacre, M. Campbell, M. Wheeler, D. Stevenson, "Flight Results from Tracking Ground Targets Using SeaScan UAVs with Gimbaling Cameras," *2007 American Control Conference*.
- [26] M. Wheeler, M. Campbell, R. Rysdyk, B. Schrick, W. Whitacre, R. Wise, "Cooperative Tracking of Moving Targets by a Team of Autonomous UAV's," *2006 Digital Avionics Aerospace Conference*.

- [27] M. Campbell, M. Wheeler, "A Vision Based Geolocation Tracking System for UAV's," *2006 AIAA Guidance, Navigation and Control Conference*.
- [28] W. Whitacre, M. Campbell, "Information-Theoretic Optimization of Periodic Orbits for Cooperative Geolocation," *AIAA Guidance, Navigation and Control Conference*, Aug 2008.
- [29] M. Campbell, M. Wheeler, "A Vision Based Geolocation Tracking System for UAV's," *AIAA Journal of Guidance, Dynamics and Controls*.
- [30] W. Whitacre, M. Campbell, "Cooperative Estimation using Mobile Sensor Nodes in the Presence of Communication Loss," submitted to the *AIAA Journal of Guidance, Control and Dynamics*.

Technology Transfer

There are several areas of notable contributions in terms of technology transfer and dissemination. First, the RoboFlag software has been used by AFRL as a simulated command and control (C2) task. Recent research at AFRL using RoboFlag includes the evaluation of a flexible delegation-type interface for the control of multiple UAVs in an adversarial, team-based environment, the influence of instant messaging (IM) on team performance and communication in an adversarial, team-based environment, and investigations of spatial and temporal cognitive load as a potential mitigator of team performance and communications. The simulator has also been used by collaborators at Caltech, Vanderbilt, and a number of other universities.

Second, the general formal, probabilistic approach to considering humans has been adopted by a number of groups, and is now a focus of research and even MURI programs. Dr. Scott Galster has transferred some of the topics to his 6.2 MIRO program at AFRL.

Finally, there are of course many publications to the community, as well as talks to universities (including underrepresented colleges), and several keynote addresses on the topic. Partially from this work, the PI achieved the Associate Fellow status at the AIAA.

Education

This program has also educated a number of undergraduate and graduate students, as well as Post-Doctoral scholars. The list below includes the graduate students and Post-Docs.

Jesse Veverka, MS 2006

David Schneider, PhD 2006

Nisar Ahmed, PhD candidate (also NSF fellowship winner)

Danelle Shah, PhD candidate (also NDSEG fellowship winner)

William Whitacre, PhD candidate

Frederic Bourgault, Post-Doc

Many undergraduates also are introduced to their first research experiences, specifically using the testbeds; a portion of this work has led to publications with undergraduate students.

Appendix: Partial Publications from the Program